The Variance Gamma Scaled Self-Decomposable Process in Actuarial Modelling

Conall O'Sullivan
University College Dublin *

Michael Moloney
Mercer IC

Initial draft: 15 May 2007
This draft: June 2010

Abstract

A scaled self-decomposable stochastic process put forward by Carr, Geman, Madan and Yor (2007) is used to model long term equity returns and options prices. This parsimonious model is compared to a number of other one-dimensional continuous time stochastic processes (models) that are commonly used in finance and the actuarial sciences. The comparisons are conducted along three dimensions: the models ability to fit monthly time series data on a number of different equity indices; the models ability to fit the tails of the times series and the models ability to calibrate to index option prices across strike price and maturities. The last criteria is becoming increasingly important given the popularity of capital gauranteed products that contain long term imbedded options that can be (at least partially) hedged by purchasing short term index options and rolling them over or purchasing longer term index options. Thus we test if the models can reproduce a typical implied volatility surface seen in the market.

Keywords: Variance gamma, regime switching lognormal, long term equity returns.
JEL Classification: G13, G23.

1   Introduction

The lognormal (LN) process is the most popular stochastic process used to model stock prices despite some serious shortcomings. It is well documented that the LN process fails to capture certain time series properties of stock prices, such as discontinuous jumps and volatility clustering. The LN process also results in returns that are normally distributed with zero skewness and zero excess kurtosis, in conflict with the distribution of most stock and index returns that exhibit significant skewness and excess kurtosis. The LN process also fails to fit the fat tails observed in the market where extreme events happen more frequently that the LN process predicts. Options markets demonstrate the pricing and hedging potential of financial models. The LN process and the associated Black-Scholes (BS) options pricing model fails to capture certain properties of the options markets. If the assumptions underlying the BS option pricing model were correct, the BS implied volatilities for options on the same underlying asset would be constant for different strike prices and maturities. However in reality BS implied volatilities are varying over strike price and maturity in what is known as the implied volatility surface. This effect comes from two sources: the data generating process is different from a LN process as evidenced by time series analysis and financial markets are incomplete whereby it is impossibe to perfectly replicate an option by dynamic trading in the underlying asset and a risk-free bond. In this paper we focus on the former and look at a number of different choices for the underlying stochastic process and test these stochastic processes in terms of their ability to fit time series data, with special emphasis on the tails of the data and the models ability to fit a range of different option price data.

Many alternative continuous time stochastic processes to LN have been proposed in the literature to address the shortcoming mentioned above. Some of the most popular have been jump-diffusion processes, Lévy processes, regime switching processes, stochastic volatility processes, and mixtures of these. They have addressed the shortcomings mentioned above with some degree of success. In this paper a parsimonious stochastic process known as the variance gamma scalable self-decomposable (VGSSD) process is compared to a number of other continuous time one-dimensional stochastic processes in terms of their ability to fit underlying time series data, tails of the time series data and in terms of their option price calibration performance. The alternative models the VGSSD is compared to include the lognormal model, a continuous time version of the regime switching lognormal model and the variance gamma model. It should be noted that all the models used are continuous

time models and the VGSSD model has no known density function so parameter estimation and derivatives pricing is carried out for all models using the characteristic function of the model. This means that the approach taken in this paper can be applied to a wide range of models with a closed form characteristic function. These include Heston's (1993) stochastic volatility model, stochastic volatility jump-diffusion models, such as Bates (1996), and a large number of Lévy processes, see Schoutens (2003), and references therein.

The remainder of the paper is organised as follows. Section 2 introduces the various models used in the paper. In section 3 the model parameters are estimated using financial time series on three indices, the S&P 500, the TSE 300 and the FTSE 100. It should be emphasised that in this section we are concerned with the real world measure, sometimes known as the P-measure. Section 4 goes on to examine the tail behaviour of the models in the P-measure. Section 5 is concerned with calibrating the models to derivatives prices on a given day using a number of different strike prices and maturities. In this case we are concerned with the risk neutral or the Q-measure. The relevance of a models options pricing ability is discussed and section 6 concludes.

2 Continuous time models for modelling long term equity returns and option prices

In this section we give a brief introduction to the various models considered in the paper and focus particular attention on the VGSSD model. Continuously compounded returns (referred to as returns when the context is clear) are denoted as $X(t) = \ln [S(t)/S(t-1)]$ where $S(t)$ is the stock (or stock index) price at time t.

2.1 Lognormal model

The lognormal (LN) process models continuously compounded returns as an arithmetica Brownian motion so that

$$X(t) = \nu t + \sigma W(t), \tag{1}$$

where $W(t)$ is a standard Wiener process, ν is the instantaneous drift and σ is the instantaneous volatility of the returns. With the use of Itô's lemma this can be formulated into

the following well known stochastic differential equation for the stock price

$$dS(t) = S(t) \left(\mu dt + \sigma dW(t) \right), \tag{2}$$

where μ is the growth rate of the stock and is related to ν as follows $\nu = \mu - \frac{1}{2}\sigma^2$. This can be integrated to yield the following formula for the dynamics of the stock or stock index price conditional on the initial stock price $S(0)$

$$S(t) = S(0) \exp \left(\left(\mu - \sigma^2/2 \right) t + \sigma W(t) \right). \tag{3}$$

As it will be used in later sections we introduce the characteristic function for the lognormal process in this section for completeness. The characteristic function is given by

$$\phi_{X(t)}(u) = E_0^P \left[e^{iuX(t)} | X(0) \right] = e^{\left(iu\nu - \frac{1}{2}u^2\sigma^2 \right)t}, \tag{4}$$

where i is the imaginary number $\sqrt{-1}$ and u is a Fourier transform variable. This model results in returns that are normally distributed and the famous Black Scholes option pricing model for derivatives.

2.2 Variance gamma model

The variance gamma (VG) process is a popular Lévy process used in financial modelling introduced by Madan and Seneta (1990), Madan and Milne (1991) and Madan, Carr and Chang (1998). The idea is to model stock price movements occurring on business time rather than on calendar time using a time transformation of a Brownian motion. The resulting model is a four parameter model where roughly speaking we can interpret the parameters as controlling the location, volatility, skewness and kurtosis of the underlying returns distribution. Closed form option pricing formulas exist under the VG model, see Madan, Carr and Chang (1998), however they involve the computation of the modified Bessel function of the second kind. Thus it is more efficient to use the Fourier transform method of Carr and Madan (1999) that utilises knowledge of the characteristic function. The gamma process is used to transform from calendar to business time. The analogy commonly used is that when the random time change speeds up the calendar clock the market is more turbulent and when the random time change slows down the calendar clock the market is more tranquil. The gamma process, like the Poisson process, is a pure jump process and this results in the VG process being a pure jump process with no diffusion component. In fact jumps of negligible size arrive infinitely often in the VG model and

3

this infinite activity allows the model to behave like a diffusion process for small jumps. Jumps of non-negligible size occur with a finite frequency and the arrival rate of these jumps decreases monotonically with the jump size. Thus the VG model can accommodate non-diffusive jumps without the use of an orthogonal Poisson jump process.

The gamma process is a subordinator, i.e. it is a stochastic process which starts at zero and has stationary and independent gamma distributed increments, see Schoutens (2003). More precisely, time enters in the first parameter: the gamma process $\gamma(t)$ follows a Gamma(at, b) law where the gamma probability density function Gamma(a, b) is given by

$$f_{\text{Gamma}}(x; a, b) = \frac{b^a x^{a-1}}{\Gamma(a)} e^{-bx}, \qquad x > 0.$$

The characteristic function of the gamma process $\gamma(t)$ is given by

$$\phi_{\gamma(t)}(u) = E\left[e^{iu\gamma(t)}\right] = \left(1 - \frac{iu}{b}\right)^{-at}.$$

The variance gamma process uses a gamma process to time change a Brownian motion. Rather than evaluate a Brownian motion at time t it is evaluated at time $\gamma(t)$ where $\gamma(t)$ follows a gamma process with $E[\gamma(t)] = t$ and var$[\gamma(t)] = \nu t$. To do this choose $a = \frac{t}{\nu}$ and $b = \frac{1}{\nu}$ so that the characteristic function of the process $\gamma(t)$ can be written as

$$\phi_{\gamma(t)}(u) = (1 - iu\nu)^{-\frac{t}{\nu}}.$$

Let $b(t; \theta, \sigma)$ denote a Brownian motion with drift

$$b(t; \theta, \sigma) = \theta t + \sigma W(t), \tag{5}$$

where θ and σ are respectively the instantaneous drift and volatility and $W(t)$ is a standard Brownian motion. From the lognormal section we know that the characteristic function of this process is given by

$$\phi_{b(t)}(u) = e^{\left(iu\theta - \frac{1}{2} u^2 \sigma^2\right)t}.$$

Madan, Carr and Chang (1998) define a VG process, $X(t; \sigma, \nu, \theta)$, as a time changed Brownian as follows

$$X(t; \sigma, \nu, \theta) = \theta\gamma(t) + \sigma W(\gamma(t)). \tag{6}$$

The density function of the VG process is known in closed form and requires the computation of the modified Bessel function of the second kind which can be time consuming. Thus

as with many Lévy processes it is sometimes more convenient to work with the characteristic function of the process which can be found by conditioning on the jump $\gamma(t)$ and is given by

$$
\begin{aligned}
\phi_{X(t)}(u) &= \left(1 - \nu\left(\mathrm{i}u\theta - \frac{1}{2}u^2\sigma^2\right)\right)^{-\frac{t}{\nu}}, \\
&= \left(1 - \mathrm{i}u\nu\theta + \frac{1}{2}u^2\nu\sigma^2\right)^{-\frac{t}{\nu}}.
\end{aligned}
\tag{7}
$$

The dynamics of the stock or stock index price are defined as

$$
S(t) = S(0)\exp((\mu + \omega)\,t + X(t; \sigma, \nu, \theta)),
\tag{8}
$$

where μ is the instantaneous expected return of the stock evaluated at calendar time and ω is a compensator term that is chosen to ensure that

$$
E_0^P\left[S(t)\right] = S(0)\exp(\mu\,t).
$$

Comparing equations 3 and 8 it can be seen that the Brownian motion volatility term in equation 3, $\sigma W(t)$, that depends on one parameter σ has been replaced with a VG random variable, $X(t; \sigma, \nu, \theta)$, that depends on three parameters, namely σ, ν and θ. The standard lognormal compensator term (also known as a convexity correction) $-\sigma^2/2$ has also been replaced with a more general compensator term ω which is easily derived from knowledge of the characteristic function (see below).

The risk neutral process used for option pricing has the following dynamics

$$
S(t) = S(0)\exp((r - q + \omega^*)\,t + X(t; \sigma^*, \nu^*, \theta^*)),
\tag{9}
$$

where r and q are the continuously compounded risk-free rate and dividend yield and the vector $\{\sigma^*, \nu^*, \theta^*\}$ contains the risk neutral parameters that need not be equal to their real-world counterparts[1] unlike in the diffusion case when the volatility parameter must be the same in both measures. A discussion on the appropriate measure change is beyond the scope of this paper. The approach taken in this paper is to imply the risk neutral parameters from a range of options prices on a given day by calibrating model option prices to market option prices and imposing that the growth rate of the stock is equal to the continuously compounded risk neutral growth rate $r - q$.

---

[1] See Madan, Carr and Chang (1998) and Cont and Tankov (2004) for more detail on this delicate issue.

The characteristic function for the logarithm of the future stock price, $\ln S(t)$, can be derived from the characteristic function for the VG process and is given by

$$\phi_{\ln S(t)}(u) = E_0^P\left[e^{iu\ln S(t)}\right] = \exp\left\{iu\left(\ln S(0) + (\mu + \omega)t\right)\right\}\phi_{X(t)}(u). \tag{10}$$

The compensator term can be found from this characteristic function and is given by $\omega = -\frac{1}{t}\ln\left(\phi_{X(t)}(-i)\right)$. This ensures the expectation of the future stock price is given by

$$E_0^P\left[S(t)\right] = E_0^P\left[e^{iu\ln S(t)}\right]\bigg|_{u=-i} = \exp\left\{\ln S(0) + \mu\,t - \ln E_0^P\left[e^{X(t)}\right]\right\}E_0^P\left[e^{X(t)}\right]$$

$$= \frac{S(0)e^{\mu t}E_0^P\left[e^{X(t)}\right]}{E_0^P\left[e^{X(t)}\right]} = S(0)e^{\mu t}.$$

The moments of the variance gamma process $X(t)$ are given by

$$E_0^P\left[X\left(t\right)\right] = \theta t, \tag{11}$$

$$\text{Var}\left[X\left(t\right)\right] = \left(\sigma^2 + \theta^2\nu\right)t, \tag{12}$$

$$\text{Skew}\left[X\left(t\right)\right] = \frac{\left(3\sigma^2\theta\nu + 2\theta^3\nu^2\right)}{\left(\sigma^2 + \theta^2\nu\right)^{\frac{3}{2}}}t^{-\frac{1}{2}}, \tag{13}$$

$$\text{Kurt}\left[X\left(t\right)\right] = 3 + 3\nu\left[2 - \frac{\sigma^4}{\left(\sigma^2 + \theta^2\nu\right)^2}\right]t^{-1}. \tag{14}$$

This results in the following real world expectation for the continuously compounded returns

$$E_0^P\left[\ln\left(\frac{S(t)}{S(0)}\right)\right] = \mu\,t - \ln\phi_{X(t)}\left(-i\right) + \theta t$$

$$= \mu\,t - \ln\left(1 - \theta\nu - \frac{1}{2}\sigma^2\nu\right)^{-\frac{t}{\nu}} + \theta t$$

$$= \left(\mu + \frac{1}{\nu}\ln\left(1 - \theta\nu - \frac{1}{2}\sigma^2\nu\right) + \theta\right)t$$

$$\to \left(\mu - \frac{1}{2}\sigma^2\right)t \text{ as } \nu \to 0.$$

It can be seen that as $\nu \to 0$ the standard lognormal convexity correction applys to the mean of the continuously compounded returns. When $\nu \neq 0$ the convexity correction is more complex. The higher moments of the continuously compounded returns are the same as the VG higher moments because the deterministic components cancel out.
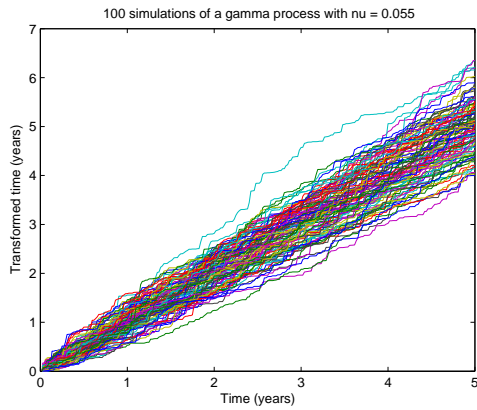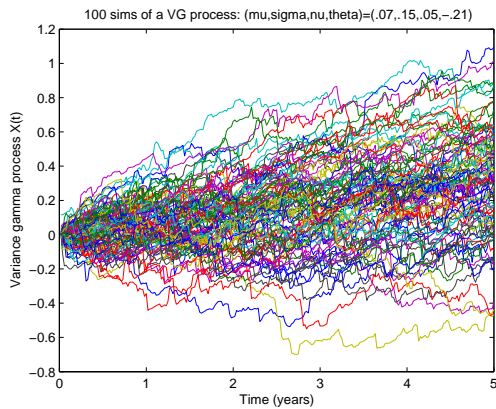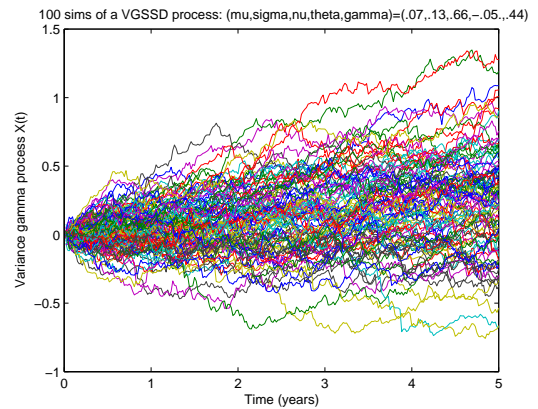
(a) $\nu = 0.0547$
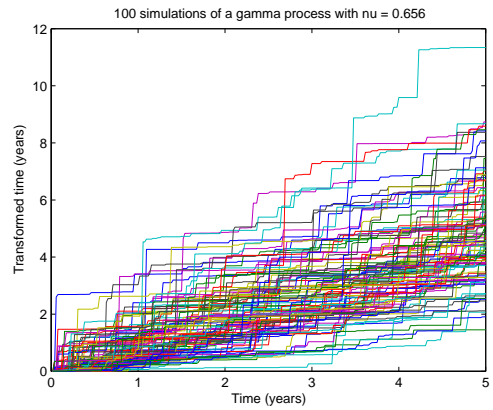
(b) $\nu = 0.6560$

Figure 1: Simulations of a gamma process

(a) $\{\mu, \sigma, \nu, \theta\} = \{0.07, 0.15, 0.05, -0.21\}$

(b) $\{\mu, \sigma, \nu, \theta, \gamma\} = \{0.07, 0.13, 0.66, -0.05, 0.44\}$

Figure 2: Simulations of a variance gamma and a variance gamma scaled self-decomposable process

2.3   Variance gamma scaled self-decomposable process

Straightforward Lévy processes such as the example described above are powerful in terms of capturing skewness and kurtosis observed in financial time series and in risk neutral density functions implied from options prices at a particular horizon. However, Lévy processes are driven by homogeneous and independent increments which fail to capture volatility clustering evident in market returns and, in a related way, do not accommodate option prices across a range of different maturities very well. Konikov and Madan (2002) show that all Lévy processes have a skewness and excess kurtosis that decreases with the length of the time horizon according to $t^{-\frac{1}{2}}$ and t^{-1} respectively (see equations 13 and 14 above for the VG case). However evidence from the options markets indicated that the higher moments implied from options prices were constant or even increasing slighty over time. Konikov and Madan (2002) proposed using a regime switching variance gamma process to model stock returns and option prices. This model is no longer a Lévy process and lacks the parsimony of the VG process as it has a total of nine parameters. On the other hand the model does provide a better fit to option prices across a wide range of strike prices and maturities and allows for two hidden regimes thus incorporating stochastic volatility by the random switching between regimes of different volatilities.

Other models have addressed this term structure of moments issue. These include stochastic volatility models proposed by Hull and White (1988) and Heston (1993) among many others. Stochastic volatility can be incorporated into a Lévy process in two ways. The first is to allow the volatility parameter to be a stochastic process and the second method is to time change a Lévy process where the second time change operates on the time t in the exponent of the characteristic function. These models require between six and ten parameters and involve a two-dimensional data generating process. This motivated Carr, Geman, Madan and Yor (2007) to consider more parsimonious models based on one-dimensional Lévy processes. Their idea was to construct stochastic processes that had inhomogeneous independent increments from Lévy processes with homogeneous independent increments. They contructed these stochastic processes in a way that rendered their higher moments constant over the maturity horizon. We only consider one of the models proposed by CGMY (2007) which is built using the variance gamma process. Their results indicated that this was one of the more successful models that they considered.

A self-decomposable random variable has the same distribution of as a scaled version of itself and an independent residual random variable. The variance gamma process is an

example of a self-decomposable process. Self-decomposable processes are Lévy processes with jump arrival rates that are decreasing in the jump size. There other specific technical constraints on the characteristic function for a Lévy process to be a self-decomposable process, see Schoutens (2003) for more information on this. A self-decomposable random variable also has a distribution of class L which means it can motivated as a limit law with more general scaling than the Gaussian limit law. This means that self-decomposable processes can be motivated as limit laws where the independent infleunces being summed are of different orders of magnitude. Thus they are appropriate building blocks for stochastic processes used to model financial markets. However self-decomposable processes (since they are a subset of Lévy processes) have higher moments that depend on the maturity horizon. This is why CGMY (2007) modelled returns using scaled self-decomposable processes.

The variance gamma scaled self-decomposable (VGSSD) stochastic process can be constructed from the variance gamma stochastic process as follows: define the scaled stochastic process $X(t)$ such that it is in equal in law to $t^\gamma X_{VG}(1)$ where $X_{VG}(1)$ is a variance gamma random variable at unit time. It follows that the characteristic function of $X(t)$ is given by

$$\phi_{X(t)}(u) = \phi_{X_{VG}(1)}(ut^\gamma) = \left(1 - iut^\gamma\nu\theta + \frac{1}{2}u^2t^{2\gamma}\nu\sigma^2\right)^{-\frac{1}{\nu}}. \tag{15}$$

VGSSD is a scaled stochastic process so its higher moments remain constant with the maturity horizon. The moments of the process are given by

$$E\left[X(t)\right] = \theta t^\gamma, \tag{16}$$

$$\text{var}\left[X(t)\right] = \left(\sigma^2 + \theta^2\nu\right)t^{2\gamma}, \tag{17}$$

$$\text{Skew}\left[X(t)\right] = \frac{\left(3\sigma^2\theta\nu + 2\theta^3\nu^2\right)}{\left(\sigma^2 + \theta^2\nu\right)^{\frac{3}{2}}}, \tag{18}$$

$$\text{Kurt}\left[X(t)\right] = 3 + 3\nu\left[2 - \frac{\sigma^4}{\left(\sigma^2 + \theta^2\nu\right)^2}\right]. \tag{19}$$

As far as the authors know the VGSSD model does not have a closed form density function thus one must use the characteristic function when evaluating the model using time series data or option prices. We now model the stock price according to equation 8 replacing the VG process $X(t; \sigma, \nu, \theta)$ with the VGSSD process $X(t; \sigma, \nu, \theta, \gamma)$. The risk neutral dynamics of the stock or index price are modelled in the same way as with the VG process where the risk neutral parameters are allowed to be different from their real world values. In both the real and risk neutral worlds the appropriate compensator terms, $\omega(t)$ and $\omega^*(t)$,

must be used. These are derived from the VGSSD characteristic function with $\omega(t) = -\frac{1}{t} \ln \phi_{X(t)}(u)$. This results in the following real world expectation for the continuously compounded returns

$$
\begin{aligned}
E_0^P \left[\ln \left(\frac{S(t)}{S(0)} \right) \right] &= \mu t - \ln \phi_{X(t)}(-\mathrm{i}) + \theta t^\gamma \\
&= \mu t - \ln \left(1 - \theta \nu t^\gamma - \frac{1}{2} \sigma^2 \nu t^{2\gamma} \right)^{-\frac{1}{\nu}} + \theta t^\gamma \\
&= \mu t + \frac{1}{\nu} \ln \left(1 - \theta \nu t^\gamma - \frac{1}{2} \sigma^2 \nu t^{2\gamma} \right) + \theta t^\gamma \\
&\to \mu t - \frac{1}{2} \sigma^2 t^{2\gamma} \text{ as } \nu \to 0 \\
&\to \left(\mu - \frac{1}{2} \sigma^2 \right) t \text{ as } \nu \to 0 \ \& \ \gamma \to \frac{1}{2}.
\end{aligned}
$$

2.4 Regime switching lognormal process

In this paper a two-state regime switching lognormal (RSLN) stochastic process, where the regimes are driven by a continuous time Markov switching process, is used as a benchmark model given its popularity in the actuarial literature. Hardy (2001) contains a very thorough review of the relevance of the regime switching lognormal model in modelling long term returns and option prices. Hardy uses a discrete time Markov switching process in her paper. This model can switch regimes from one interval to the next but not in between intervals. In this study to remain consistent with the other models used in the paper a regime switching lognormal process with a continuous time Markov switching process is used. Guo (2001) introduced such a model and derived option pricing formula in terms of an integral of a Bessel function. This semi-closed form solution is time consuming for the purposes of estimation and calibration. For reasons of computational speed, and to remain consistent with the other models used in this paper, it is preferred to work with the characteristic function of the stochastic process rather than the density function of the process or the known option price formulae. Konikov and Madan (2002) derived the characteristic function for a regime switching variance gamma process where the regime switch follows a continuous time Markov switching process. In this paper this characteristic function is adapted to the regime switching lognormal case.

The two-state RSLN model assumes the following dynamics for the returns

$$X\left(t\right) = \int_0^t \left[\left(1 - U\left(s\right)\right)\mathrm{d}X_1\left(s\right) + U\left(s\right)\mathrm{d}X_0\left(s\right)\right], \tag{20}$$

where $X_0\left(t\right)$ and $X_1\left(t\right)$ are lognormal processes with

$$X_0\left(t\right) = \left(\mu_0 - \sigma_0^2/2\right)t + \sigma_0 W_0(t),$$
$$X_1\left(t\right) = \left(\mu_1 - \sigma_1^2/2\right)t + \sigma_1 W_1(t),$$

and where $U(t)$ is a two-state Markov chain that takes values in the set $\{0,1\}$ with state transition rates given by parameters λ_{01} and λ_{10}. The probability that the current state is regime 0 is given by the parameter p. Denoting the characteristic functions of the individual lognormal processes at unit time $(t = 1)$ as ϕ_0 and ϕ_1 (see equation 4), the characteristic function of the regime switching lognormal process is given by

$$\phi_{X(t)}\left(u\right) = \phi_0\left(u\right)^t g\left(\ln\left(\frac{\phi_0\left(u\right)}{\phi_1\left(u\right)}\right)\right), \tag{21}$$

where

$$g\left(\lambda\right) = pg_0\left(\lambda\right) + (1-p)g_1\left(\lambda\right),$$

$$g_0\left(\lambda\right) = e^{-(\eta_1(\lambda)+\lambda_{01})t} \times \frac{\eta_2\left(\lambda\right) + \lambda_{01} - \left(\eta_1\left(\lambda\right) + \lambda_{01}\right)e^{-(\eta_2(\lambda)-\eta_1(\lambda))t}}{\eta_2\left(\lambda\right) - \eta_1\left(\lambda\right)},$$

$$g_1\left(\lambda\right) = (1/\lambda_{01})e^{-(\eta_1(\lambda)+\lambda_{01})t} \times \frac{\eta_2\left(\lambda\right)\left(\eta_1\left(\lambda\right) + \lambda_{01}\right)e^{-(\eta_2(\lambda)-\eta_1(\lambda))t} - \eta_1\left(\lambda\right)\left(\eta_2\left(\lambda\right) + \lambda_{01}\right)}{\eta_2\left(\lambda\right) - \eta_1\left(\lambda\right)},$$

$$\eta_1\left(\lambda\right) = \frac{\lambda + \lambda_{10} - \lambda_{01}}{2} - \sqrt{\frac{(\lambda + \lambda_{10} - \lambda_{01})^2}{4} + \lambda_{10}\lambda_{01}},$$

$$\eta_2\left(\lambda\right) = \frac{\lambda + \lambda_{10} - \lambda_{01}}{2} + \sqrt{\frac{(\lambda + \lambda_{10} - \lambda_{01})^2}{4} + \lambda_{10}\lambda_{01}}.$$

This characteristic function is derived by recognising that the Laplace transform of the time spent in regime 1 is known in analyical form. A detailed derivation is beyond the scope of this paper however for more details on this derivation see Konikov and Madan (2002).

Denote τ_{ij} as the time that the regime switches from state i to state j. Given that the current state is i, the probability of remaining in state i and not switching to state j over the time period $(0, t)$ is given by

$$\Pr\left\{\tau_{ij} > t\right\} = \exp(-\lambda_{ij}t), \text{ for } i, j \in \{0, 1\} \text{ and } j \neq i.$$

To reduce the number of parameters in the model we assume that the probability of switching states is equal to one minus the probability of remaining in the current state

$$p_{ij} = \Pr\{\tau_{ij} < t\} = 1 - \exp(-\lambda_{ij}t), \text{ for } i, j \in \{0, 1\} \text{ and } j \neq i,$$

and then use Hardy (2001) to write p (the unconditional probability of being in state 0) in terms of p_{01} and p_{10} with $p = p_{01}/(p_{01} + p_{10})$.

The dynamics of the stock or stock index price are defined by

$$S(t) = S(0)\exp\left(X\left(t; \mu_0, \sigma_0, \mu_1, \sigma_1, \lambda_{01}, \lambda_{10}\right)\right). \tag{22}$$

It should be noted that unlike the other models used in this paper the growth rate of the stock in the real world measure is not explicitly modelled as μ but is a function of the model parameters $\mu_0, \sigma_0, \mu_1, \sigma_1, \lambda_{01}$, and λ_{10}. This is to ensure that the model is comparable to other regime switching lognormal processes used in the literature such as Hardy (2001). The mean of the RSLN process can be derived from the characteristic function with $E[X(t)] = \frac{1}{i}\frac{\partial\phi_{X(t)}(u)}{\partial u}|_{u=0}$.[2] The risk neutral dynamics of the stock or stock index price are defined by

$$S(t) = S(0)\exp\left((r - q + \omega^*(t))t + X\left(t; \mu_0^*, \sigma_0^*, \mu_1^*, \sigma_1^*, \lambda_{01}^*, \lambda_{10}^*\right)\right), \tag{23}$$

where in this paper we imply the risk neutral parameters from market option prices and allow them to be different from their real world counterparts. By using the above form for the risk neutral process a constraint is implicitly imposed on the risk neutral parameter vector $\{\mu_0^*, \sigma_0^*, \mu_1^*, \sigma_1^*, \lambda_{01}^*, \lambda_{10}^*\}$ so that the risk neutral growth rate of the stock price is $r - q$ i.e. $E_0^Q[S(t)] = S(0)e^{(r-q)t}$.

3   Time series data and estimation methodology

In this section the LN, VG and VGSSD and RSLN models are are estimated using monthly total returns data on the TSE 300, the S&P 500 and the FTSE 100. The data on the TSE 300 and the S&P 500 span the dates from 31/01/1956 to 31/12/1999 so that results are comparable to Hardy's (2001) results. The models are also estimated using FTSE 100 total returns from 31/01/1986 to 29/12/2006. The parameters of the models are esimated

---

[2]In fact all the moments can be derived from the knowledge of the characteristic function since $E[X(t)^n] = \frac{1}{i^n}\frac{\partial^n\phi_{X(t)}(u)}{\partial u^n}|_{u=0}$. However the moments of the RSLN are not reported as they are take up too much space.

using an approximate maximum likelihood estimation (MLE) method similar to the one used by Carr, Geman, Madan and Yor (2002). For a given parameter vector the density function of the stochastic process is calculated at N points $y_1, y_2, \ldots, y_N$ (where $N = 2^{14}$) over a finite range by inverting the characteristic function with the use of a fast Fourier transform (FFT)[3]. Given m observed data points CGMY (2002) arrange this observed data x_i for $i = 1, \ldots, m$ into their corresponding intervals $x_i \in [y_j, y_{j+1}]$ for $j = 1, \ldots, N-1$ and count the number of observed data points that fall into each interval (in many cases this is zero). The likelihood of observing this binned data is then maximised by appropriate choice of the parameter vector. This method involves a form of smoothing where a histogram of the data is evaluated. However rather than binning the observed data the approach taken in this paper is to evaluate the density function at the observed data points, $f(x_i)$ for $i = 1, \ldots, m$, by interpolation where $f(x_i)$ is interpolated using its $2k+1$ nearest neighbours: $f(y_{j-k}), f(y_{j-k+1}), \ldots, f(y_{j+k})$. The loglikelihood function of this interpolated density function is then maximised by appropriate choice of the parameter vector. This method can be compared to the standard maximum likelihood approach (that uses the closed form density function) however the approximate method introduces interpolation error and the loglikelihood values from the approximate method and the standard MLE method will not be exactly the same.

The parameter values of the LN, VG, VGSSD and the RSLN models for the TSE 300, the S&P 500 and the FTSE 100 data are shown in Table 1. The LN, VG and VGSSD models have similar instantaneous drift and volatility parameters in all cases. The volatility parameter in the VG model is always a little lower than the volatility parameter in the LN model. This is because the standard deviation of the VG process is attributable to the three parameters $\{\sigma, \nu, \theta\}$. In the VG model the parameters ν and θ are similar across the different markets with $\nu \approx 0.03$ to 0.05 and $\theta \approx -0.17$ to -0.20. The VGSSD model has a larger ν parameter and a smaller θ parameter than the VG model. This is because the VGSSD random variable at time t is equivalent to a scaled VG random variable at unit time, thus to induce similar levels of skewness and kurtosis as those in the VG model the moments in the VGSSD model need to have a higher ν and a lower θ. The γ parameter in the VGSSD model is always close to 0.5 which is what one expects if markets are very nearly efficient since the variance of the returns grow proportional to time t when $\gamma = 0.5$. The parameter values of the RSLN model are very different from the other three models

---

[3]See appendix on how to invert the characteristic function to obtain the density function.

LN,VG and VGSSD	Model parameters					
	μ	σ	ν	θ	γ	
RSLN	μ_0	σ_0	μ_1	σ_1	λ_{01}	λ_{10}
TSE 300 (1956 - 99 Monthly Total Returns)						
LN	0.0610	0.1561	-	-	-	
VG	0.0605	0.1482	0.0409	-0.2005	-	
VGSSD	0.0605	0.1685	0.4912	-0.0658	0.5516	
RSLN	0.1708	0.1199	-1.2320	0.4372	2.0771	74.5509
S&P 500 (1956 - 99 Monthly Total Returns)						
LN	0.0695	0.1436	-	-	-	
VG	0.0697	0.1376	0.0340	-0.1767	-	
VGSSD	0.0692	0.1463	0.3767	-0.0672	0.5271	
RSLN	0.2341	0.0842	-0.2212	0.2558	12.8848	61.8610
FTSE 100 (1986 - 2006 Monthly Total Returns)						
LN	0.0673	0.1609	-	-	-	
VG	0.0666	0.1493	0.0547	-0.2072	-	
VGSSD	0.0666	0.1301	0.6560	-0.0521	0.4447	
RSLN	0.1488	0.1370	-0.0149	0.3948	0.0203	0.3974

Table 1: Comparison of parameters for the lognormal, variance gamma, variance gamma scaled self-decomposable and regime switching lognormal processes.

which is not surprising given that the RSLN model is based a different paradigm than the other three models. The results for the RSLN model can be interpreted in a consistent manner across the different markets. There is a low volatility regime with a positive drift and a high volatility regime with a negative drift and the process switches out of the high volatility regime very quickly relative to the low volatility regime. What is surprising is the results for the TSE data where the high volatility regime has a very large negative return of -123.2%! However the probability of remaining in this regime for a length of period t is equal to $e^{-74.5509t} = 0.2\%$ for $t = 1/12$. Thus although this regime has a very large negative drift the probability of switching out of this regime is very large.

The maximum likelihood results for the three markets: TSE 300, S&P 500 and FTSE 100 are shown in Table 2. Similar to Hardy (2001) the following results are reported: the

log likelihood function (LL), the Schwartz-Bayes information criteria (SBC), the Akaike information criteria (AIC) and the likelihood ratio test (LRT) versus the LN model. The loglikelihood method selects the model with the maximum value for LL. In the interests of parsimony the AIC selects the model with the maximum value for $LL - n$ where n is the number of parameters in the model. This captures in an ad-hoc fashion the fact that each additional parameter of the model should contribute at least one unit to the loglikelihood value. The SBC selects the model with the maximum value for $LL - \frac{1}{2}n \ln m$, where n is the number of parameters and m is the number of observed data points with $m = 527$ for the TSE 300 and the S&P 500 and $m = 252$ for the FTSE 100. For a sample size of 527 (252) each additional parameter must increase the loglikelihood value by at least 3.13 (2.76). This is a more formal information criteria than the AIC and puts more weight on parsimonious models than the AIC. The likelihood ratio compares embedded models where a model with n_1 parameters is a special case of a model with n_2 parameters where $n_2 > n_1$. Under the null hypothesis that there is no improvement under model 2 the test statistic $2(LL_2 - LL_1)$ has a χ^2 distribution with degrees of freedom equal to $n_2 - n_1$. In this paper the LN is a special case of the VG model (when $\nu \to 0$), but the LN and VG are not special cases of the VGSSD model. Also the LN model is a special case of the RSLN. However even for models that are not imbedded the likelihood ratio test can still be used for model selection although the χ^2 is an approximation for the true distribution of the test statistic.

It is clear that all three of the VG, VGSSD and RSLN models provide a better fit according to all three measures (Log-likelihood, Schwartz-Bayesian Criterion and Akaike Information Criterion). VG and VGSSD models are a marginally better fit to S&P 500 and FTSE 100 data on the Log-likelihood measure and the benefit of their greater parsimony is clearly relative to the RSLN model from the SBC and AIC criteria.

4   Tail Behaviour of Models

Risks inherent in particular investment strategies will often be assessed by looking at percentiles of outcomes - either absolute in the sense of a distribution of asset returns per se or, more likely, the distribution of asset returns relative to movement in underlying liabilities - which effectively amount to use of a 'Value-at-Risk' metric. This quantile measure of risk has many problems which are well documented (see for example Artzner *et al* 1999) and awareness of these shortcomings often leads to use of the Conditional Tail Expectation

Model	Number of parameters	LL	SBC	AIC	LRT (p-value)
	TSE 300 (1956 - 99 Monthly Total Returns)				
LN	2	885.64	879.37	883.64	
VG	4	912.67	900.13	908.67	5.32e-009
VGSSD	5	912.67	897.00	907.67	5.32e-009
RSLN	6	914.23	895.43	908.23	5.32e-009
	S&P 500 (1956 - 99 Monthly Total Returns)				
LN	2	929.76	923.49	927.76	
VG	4	948.92	936.39	944.92	5.32e-009
VGSSD	5	948.92	933.26	943.92	5.32e-009
RSLN	6	948.54	929.74	942.53	5.32e-009
	FTSE 100 (1986 - 2006 Monthly Total Returns)				
LN	2	415.93	410.40	413.93	
VG	4	436.39	425.33	432.39	5.32e-009
VGSSD	5	436.39	422.57	431.39	5.32e-009
RSLN	6	432.09	415.50	426.09	5.32e-009

Table 2: Comparison of selection information for the lognormal, variance gamma, variance gamma scaled self-decomposable and regime switching lognormal processes.

| | *p*-value | | | |
Model	10	5	2.5	1
TSE 300 (1956 - 99 Monthly Total Returns)				
LN	-6.02	-38.63	-44.88	-38.73
VG	7.99	-16.71	-20.87	-16.8
VGSSD	7.99	-16.71	-20.87	-16.8
RSLN	3.52	-18.94	-19.93	-13.68
S&P 500 (1956 - 99 Monthly Total Returns)				
LN	7.62	-25.87	-33.15	-28.66
VG	15.37	-11.17	-16.5	-13.49
VGSSD	16.69	-10.27	-15.95	-13.24
RSLN	12.62	-12.48	-17.2	-14
FTSE 100 (1986 - 2006 Monthly Total Returns)				
LN	-9.03	-22.14	-24.46	-22.76
VG	0.79	-8.11	-9.26	-8.05
VGSSD	0.78	-8.11	-9.26	-8.05
RSLN	-3.94	-13.16	-11.85	-8.39

Table 3: Comparison of loglikelihood fit of the lognormal, variance gamma, variance gamma scaled self-decomposable and regime switching lognormal processes for tail distribution of data.

(CTE) measure, defined as the expected value of the loss given that the loss falls beyond a specified quantile of the distribution.

Given the importance of VaR and CTE for decision-making purposes, it is interesting to compare how well the tail of each model fits observed data and to examine differences in implied VaR and CTE over different time horizons.

As a first step in trying to gauge the overall goodness of fit of each distribution to the tail of the observed time series data, Table 3 repeats the Log-likelihood test for each distribution but in this case the likelihood function is summed across only those observations falling within the percentile shown.

The VG, VGSSD and RSLN models are a much better fit to the tail of the observed data at all significance levels. At the 10% and 5% levels for TSE data, both VG and VGSSD

Model	p-value			
	10	5	2.5	1
TSE				
LN	-0.0484	-0.0639	-0.0771	-0.0922
VG	-0.0466	-0.0677	-0.0876	-0.1125
VGSSD	-0.0466	-0.0677	-0.0876	-0.1125
RSLN	-0.0431	-0.0633	-0.0869	-0.1224
S&P 500				
LN	-0.0424	-0.0567	-0.0689	-0.0829
VG	-0.0406	-0.0594	-0.0769	-0.0988
VGSSD	-0.0414	-0.0603	-0.078	-0.1
RSLN	-0.0394	-0.0585	-0.0764	-0.0984
FTSE 100				
LN	-0.0492	-0.0651	-0.0787	-0.0943
VG	-0.0466	-0.0697	-0.0917	-0.1195
VGSSD	-0.0466	-0.0697	-0.0917	-0.1195
RSLN	-0.0417	-0.058	-0.0745	-0.1047

Table 4: Comparison of quantile risk measures of the lognormal, variance gamma, variance gamma scaled self-decomposable and regime switching lognormal processes.

models provide the best fit while at the 2.5% and 1% levels the RSLN model is ahead. For S&P 500 and FTSE100 data, the VG and VGSSD models provide the best fit across all all four tail sections.

The impact of the better fit to observed tail data for the VG, VGSSD and RSLN models is clear from the Value-at-Risk figures in Table 4 and Conditional Tail Expectation (CTE) figures in Table 5 where risk exposures are all materially understated by a LN assumption.

Figures 3 to 5 show the segment of the CDF for each fitted distribution up to the 10th percentile. In all cases the CDF has been calculated numerically from the PDF derived as the inverse Fourier transform of the characteristic function for each distribution. The parameters being considered were fit based on monthly return data and used to derive a corresponding CDF over periods ranging from 1 month to 10 years.

While the VG starts with an identical fit to the VGSSD at one month, as expected, both

	p-value			
Model	10	5	2.5	1
TSE 300 (1956 - 99 Monthly Total Returns)				
LN	-0.0683	-0.0811	-0.0924	-0.1057
VG	-0.0756	-0.0953	-0.114	-0.1376
VGSSD	-0.0756	-0.0953	-0.114	-0.1376
RSLN	-0.0754	-0.0989	-0.1247	-0.1589
S&P 500 (1956 - 99 Monthly Total Returns)				
LN	-0.0608	-0.0727	-0.0832	-0.0956
VG	-0.0663	-0.0837	-0.1001	-0.1209
VGSSD	-0.0673	-0.0847	-0.1012	-0.1221
RSLN	-0.0655	-0.0828	-0.0993	-0.1195
FTSE 100 (1986 - 2006 Monthly Total Returns)				
LN	-0.0697	-0.0828	-0.0945	-0.1081
VG	-0.0787	-0.1004	-0.1212	-0.1479
VGSSD	-0.0787	-0.1004	-0.1212	-0.1479
RSLN	-0.0689	-0.089	-0.1124	-0.1529

Table 5: Comparison of conditional tail expectations of the lognormal, variance gamma, variance gamma scaled self-decomposable and regime switching lognormal processes.
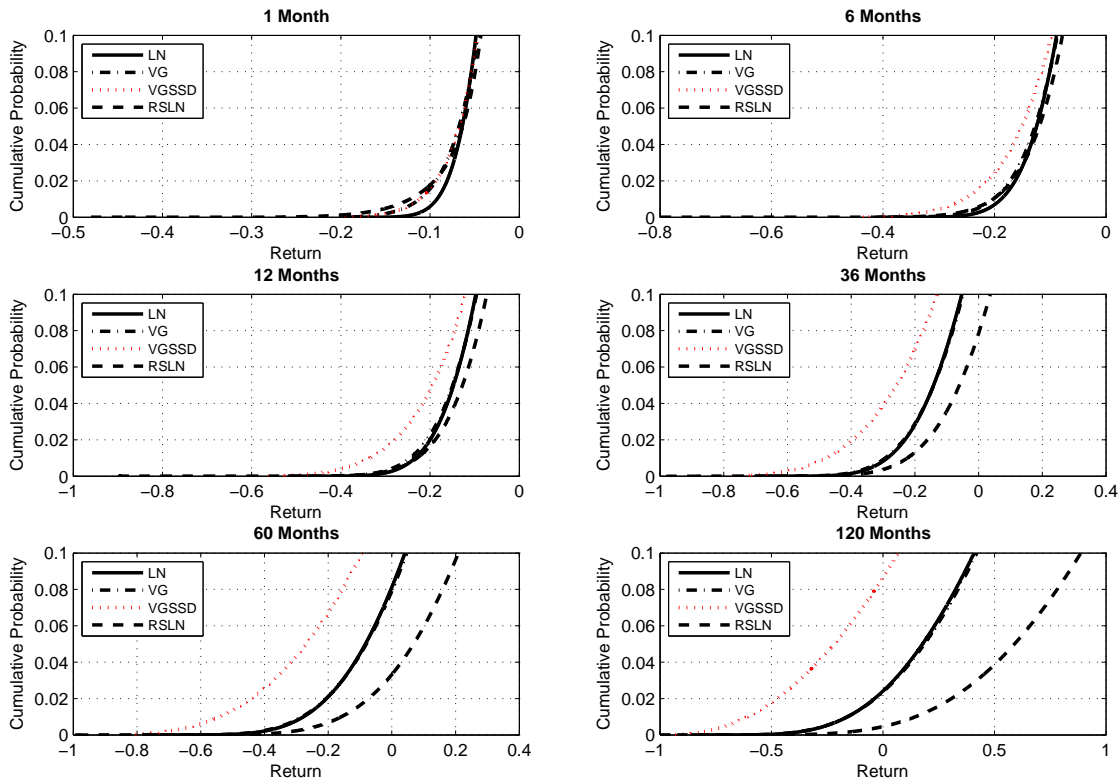
Figure 3: Implied tail distribution for TSE fit over various time horizons

its negative skewness and excess kurtosis are seen to decrease rapidly across all three fits such that from 36 months onwards it is hard to discern any difference between the graphs for VG and LN. As expected, the VGSSD model preserves both skewness and kurtosis independent of time horizon.

5   Option Pricing

In this section the models are measured on their ability to reproduce a large number of options across different strike prices and maturities on a given day. All the models used in this paper are one-dimensional Markov models. It is not possible to reproduce the dynamics
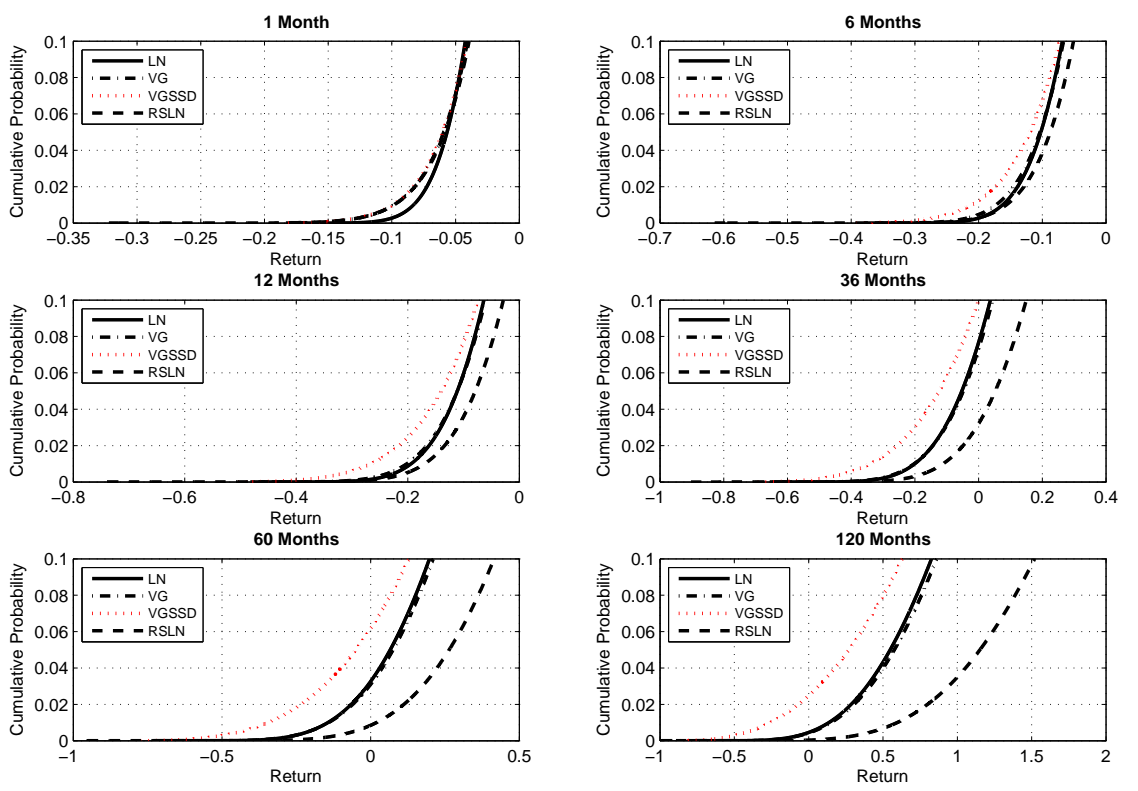
Figure 4: Implied tail distribution for S&P fit over various time horizons
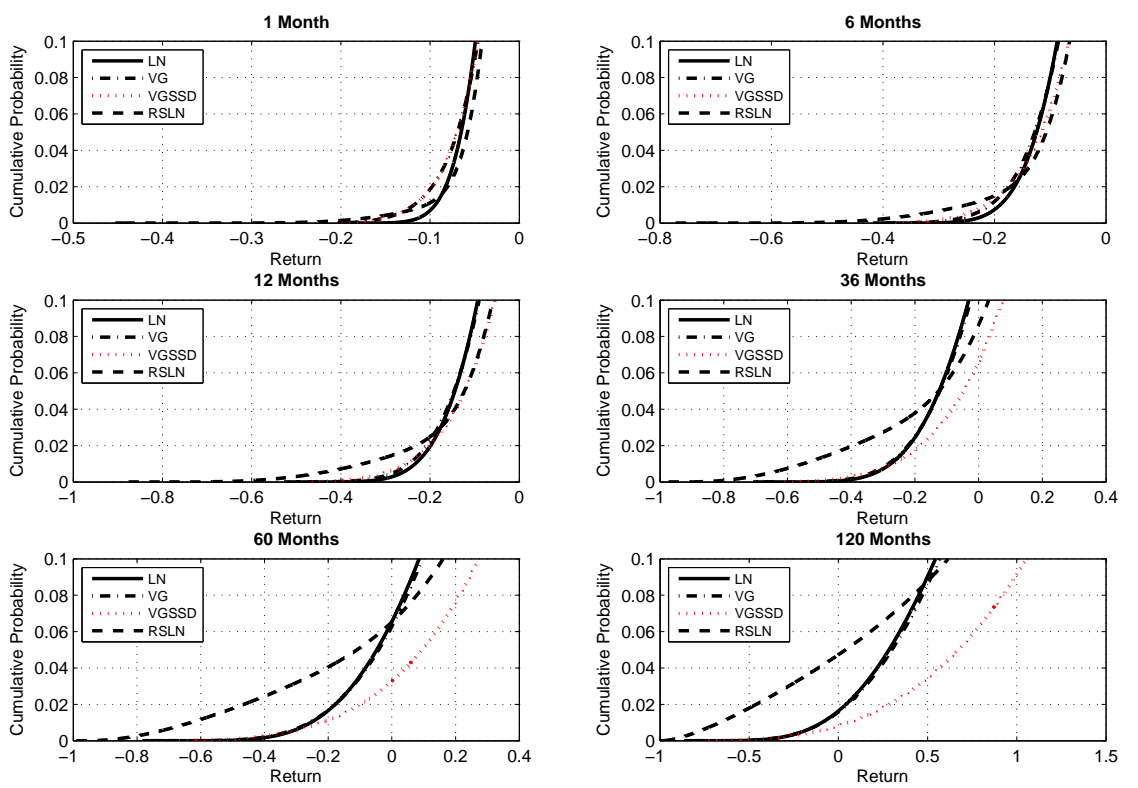
21

Figure 5: Implied tail distribution for FTSE fit over various time horizons
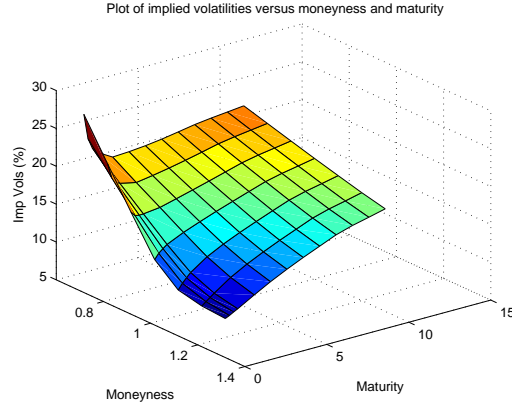
Figure 6: Plot of implied volatilities on the FTSE 100 index on the 11th January 2007.

of the implied volatility surface[4] using such models as the data generating process. However it may be possible to reproduce the average shape of the implied volatility surface over time[5] which may be useful for actuarial modelling when one is dealing with long time periods and still needs to model certain option properties such as, for example, the fact that out-of-the money put options are more expensive, relative to Black-Scholes model, than out-of-the-money call options. Models that can reproduce the average shape of the implied volatility surface may prove useful in simulating future scenarios for the underlying asset and using the simulated future underlying asset price and the risk neutral parameters of the model to generate a realistic set of future option prices written on the underlying asset. Rather than test a models ability to reproduce an average implied volatility surface we test the models ability to reproduce the implied volatility surface on a given day. This is a more difficult test because the average surface will be smoother than the surface on a given day. Figure 6 depicts the implied volatility surface on the 11th January 2007 for the FTSE 100 index.

With the exception of the lognormal model the other models used in this paper result in an incomplete market where an individual option cannot be replicated by dynamic hedging in the underlying asset and a risk-free bond. A partial equilibrium approach is followed in this paper where the price of the option is determined relative to the underlying asset and

---

[4]The grid of implied volatilities plotted against strike price and maturity.

[5]This is where we take the observed implied volatility surface over a period of time, interpolate it so we observe implied volatilities on a fixed grid of maturity and moneyness (strike/spot) and then average these implied volatilities over different dates.

it is assumed that the growth rate of the asset in the risk neutral world is equal to the risk neutral growth rate $r-q$. The risk neutral parameters are then implied from market option prices and are allowed to differ from their real-world counterparts. Thus the risk neutral parameters will reflect the risk premia implicit in option prices that is partly caused by the lack of hedging perfection.

The option pricing method used is the fast Fourier transform method of Carr and Madan (1999). This method only needs knowledge of the characteristic function and it returns option prices at a range of different strike prices with one application of a FFT using the following formula

$$c\left(K,T\right) = \frac{\exp\left(-\alpha\ln\left(K\right)\right)}{\pi}\int_{0}^{+\infty}\exp\left(-\mathrm{i}v\ln\left(K\right)\right)\psi\left(v\right)\mathrm{d}v \tag{24}$$

where

$$\psi\left(v\right) = \frac{\exp\left(-rT\right)\phi_{\ln S(T)}\left(v-\left(\alpha+1\right)\mathrm{i}\right)}{\alpha^2+\alpha-v^2+\mathrm{i}\left(2\alpha+1\right)v} \tag{25}$$

and where $\phi_{\ln S(T)}$ is the model specific characteristic function of the log stock price such as that given in equation 10 for the VG process but with the real world growth μ replaced by the risk neutral growth rate $r-q$.

5.1 Data

FTSE 100 index futures options are used to measure the calibration performance of the models in the paper. The data consists of a range of options at different strike prices and maturities on the 11th January 2007. See Table 6 for more details on the data.

The models are calibrated to this implied volatility surface data. A number of filters are run on the market option prices to be used in the calibration. Option prices that are less than $0.00075\cdot S$, where S is the underlying price, are discarded, and options with maturities less than 15 days are also discarded due to these options being less liquid. Put prices were used for strike prices less than the underlying price (for $K < S$), and call prices were used for strike prices greater than the underlying price (for $K > S$). Thus we always used out-of-the money options in the calibration. Madan, Carr and Chang (1998) show that the maximum likelihood estimates of the risk neutral parameters is asymtotically equivalent to

| | | | Strike Price | | | | |
Maturity	4361.1	4984.1	5607.1	6230.1	6853.1	7476.1	8099.1
0.2464		23.1075	18.1155	13.2689	10.6758		
0.4956	24.8123	21.1118	17.5112	13.9434	11.4142		
0.7474	23.4625	20.3828	17.4050	14.4166	12.0056	10.8380	
0.9993	22.7376	20.0367	17.4393	14.8219	12.5523	11.2521	
2.0041	21.4760	19.5455	17.6994	15.8239	14.0378	12.7362	11.8584
3.0007	21.5998	19.8901	18.2535	16.6540	15.1302	13.9227	13.0156
4.0000	21.8114	20.2893	18.8318	17.4236	16.0754	14.9535	14.0707
4.9993	22.0970	20.7352	19.4266	18.1627	16.9454	15.8935	15.0431
6.0014	22.4530	21.2231	20.0342	18.8830	17.7701	16.7770	15.9596
7.0062	22.8335	21.7255	20.6456	19.5953	18.5770	17.6475	16.8707
8.0027	23.1125	22.0918	21.0888	20.1088	19.1550	18.2681	17.5176
8.9993	23.3319	22.3784	21.4342	20.5075	19.6024	18.7476	18.0167
10.0014	23.5652	22.6761	21.7918	20.9215	20.0696	19.2568	18.5527

Table 6: Market BS implied volatilities (%) for FTSE 100 index options on the 11 January 2007. The strike prices and maturities (in years) are given in the table and the other observable inputs are $S = 6230.1, r = 0.0521$ and $q = 0.0306$.

minimising the following objective function

$$f = \sqrt{\frac{1}{N} \sum_{1}^{N} \left(\ln C_i - \ln C_i\left(\Theta\right)\right)^2},$$

where C_i is the observed market price on the i-th option and $C_i\left(\Theta\right)$ is the model price of the i-th option with parameter vector Θ. However this approach seems to put a lot of emphasis on out-of-the money options at the expensive of fitting at-the-money options. Thus in this study we minimise the average absolute percentage error (AAPE) which is given by

$$f = \frac{1}{N} \sum_{1}^{N} \left|\frac{C_i - C_i\left(\Theta\right)}{C_i}\right|.$$

In fact when the market and model prices are very close $\ln C_i - \ln C_i\left(\Theta\right) \approx \frac{C_i - C_i(\Theta)}{C_i}$ and these objective functions are very similar. Table 7 contains estimates of the risk neutral parameters as of the 11th January 2007, along with the calibration performance of each model given by the AAPE in the last column. Figures 7 - 9 depict the calibration performance of the VG, VGSSD and RSLN models by graphing model implied volatilities and market implied volatilities for a range of different moneyness levels (strike price/underlying price) and maturities. As can be seen both the VGSSD and RSLN models have the best performance with the RSLN being slighly better but at the expense of two more parameters. The calibration performance of these two models is very good given the range of different option prices that are tested. This evidence suggests that the VGSSD and RSLN models seem reasonable good models to use if one requires the model to be able to reproduce the shape of the implied volatility surface in a reliable manner.

6   Conclusion

A number of different stochastic processes suitable for long term modelling of underlying asset prices and option prices are tested in the paper using knowledge of the characteristic function. Based on evidence from time series data, in particular the tails of the data, and evidence from options prices the VGSSD and the RSLN models seem to do reasonably well on all tests. The RSLN model is a well known model in the actuarial literature, but perhaps the use of a continuous time Markov process to drive the switching process is less common and this is introduced in this paper. The VGSSD model is a more recent model that is less

	Risk neutral model parameters						AAPE (%)
LN,VG and VGSSD	μ^*	σ^*	ν^*	θ^*	γ^*		
RSLN	μ_0^*	σ_0^*	μ_1^*	σ_1^*	λ_{01}^*	λ_{10}^*	
LN	-	0.1495	-	-	-		45.54
VG	-	0.1205	0.6870	-0.1439	-		20.89
VGSSD	-	0.1182	0.5668	-0.1628	0.6372		8.93
RSLN	0.1634	0.0832	0.0090	0.2141	0.0525	0.1364	8.73

Table 7: Comparison of risk neutral parameters for the lognormal, variance gamma, variance gamma scaled self-decomposable and regime switching lognormal processes based on calibration to FTSE 100 Index options on the 11th January 2007.
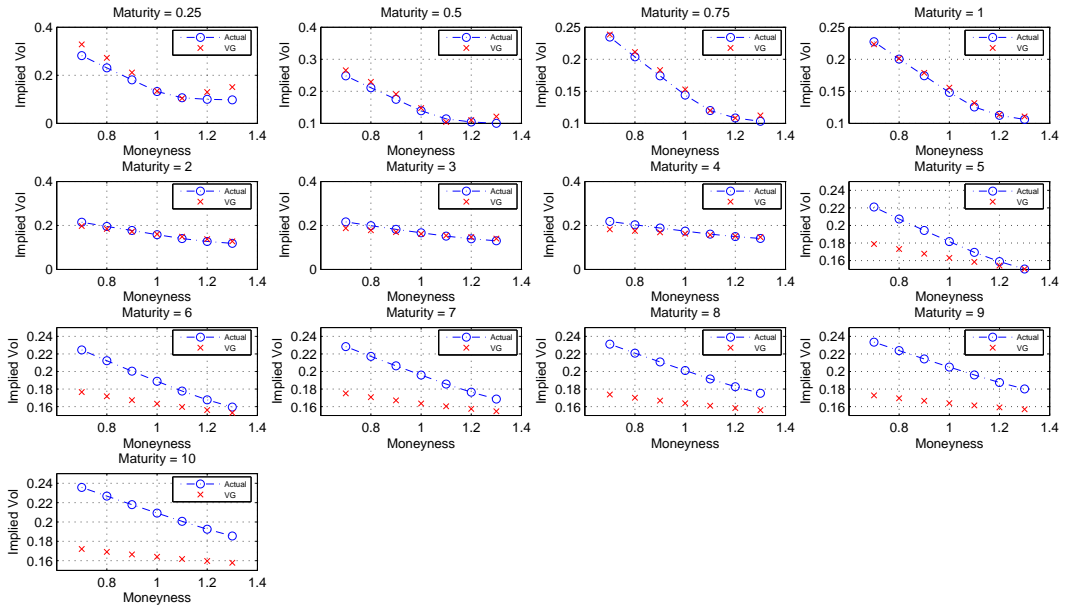

Figure 7: Market and VG implied volatilities versus moneyness for a number of different option maturities on 11th Jan 2007 for FTSE 100 index options.
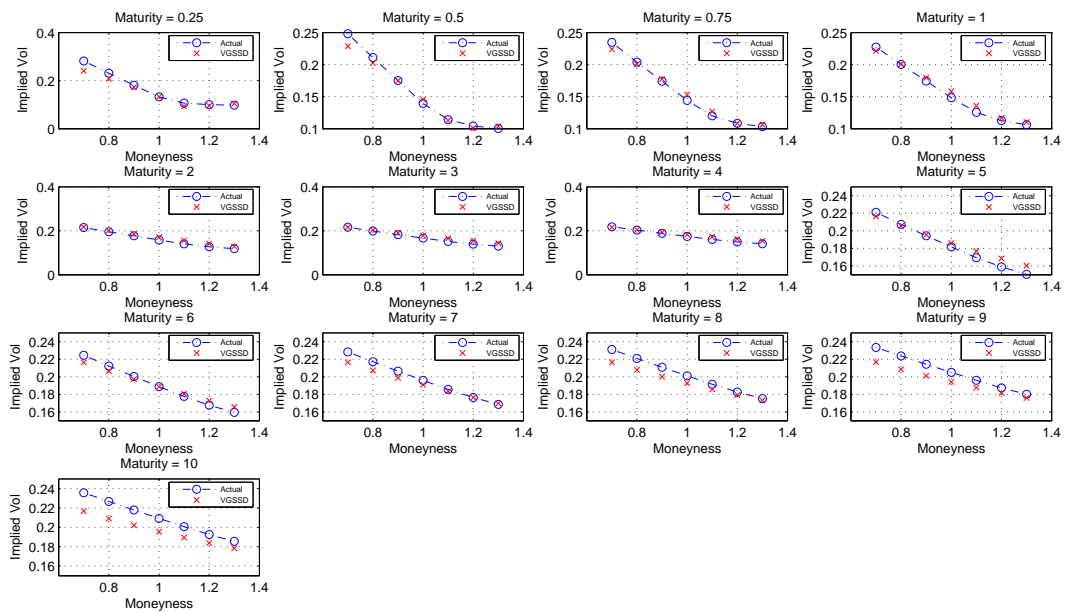
Figure 8: Market and VGSSD implied volatilities versus moneyness for a number of different option maturities on 11th Jan 2007 for FTSE 100 index options.
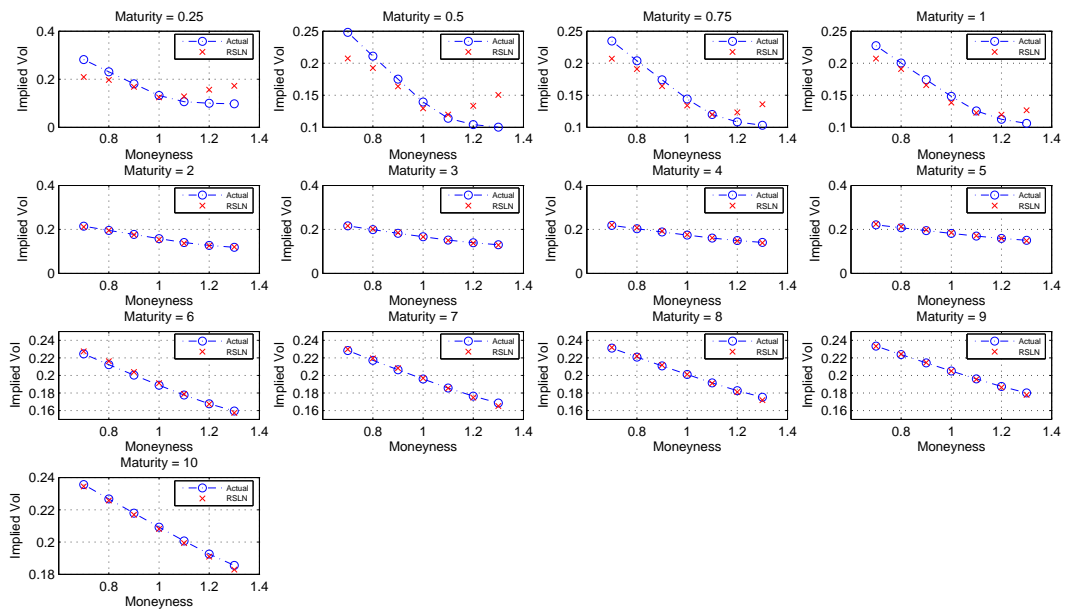
Figure 9: Market and RSLN implied volatilities versus moneyness for a number of different option maturities on 11th Jan 2007 for FTSE 100 index options.

well known and seems to perform just as good as the RSLN model however further more detailed testing is needed before more rigorous conclusions can be reached. Further research includes the analysis of options data over different time periods and different markets and the inclusion of more detailed time series and tail tests.

References

[1] Artzner, P., F. Delbaen, J.-M. Eber, and Heath, D. (1999) *Coherent Measures of Risk*. Mathematical Finance 9: 20328.

[2] Bates, D. (1996). *Jumps and stochastic volatility: the exchange rate processes implicit in Deutschemark options*. Rev. Fin. Studies, 9:69-107.

[4] Carr, P., H. Geman, D. Madan and Yor, M. (2002). *The fine structure of asset returns: an empirical investigation*. Journal of Business, 75.

[4] Carr, P., H. Geman, D. Madan and Yor, M. (2007). *Self decomposability and options pricing*. Mathematical Finance, 17 (1).

[5] Carr, P. and Madan, D.P. (1999). *Option valuation using the fast Fourier transform*, Journal of Business, 75:305-332.

[6] Cont, R. and Tankov, P. (2003). *Financial Modelling with Jump Processes*, Chapman and Hall/CRC.

[7] Cont, R. (2001). *Empirical Properties of asset returns: stylised facts and statistical issues*, Quantitative Finance (1) (2001), 223 - 236.

[8] Guo. X. (2001). *Information and option pricings*, Quantitative Finance, 1(1):38-44.

[9] Hardy, M. (2001). *A Regime Switching Model of Long-Term Stock Returns*, North American Actuarial Journal 5(2):41-43.

[10] Konikov, M. and Madan, D.B. (2002). *Option Pricing Using Variance Gamma Markov Chains*, Review of Derivatives Research, 5, 81-115.

[11] Madan, D.B., Carr,. P.P. and Chang, E.C. (1998). *The Variance Gamma Process and Option Pricing*, European Finance Review 2:79-105.

[13] Madan, D.B. and Milne, F. (1991). *Options pricing with variance gamma martingale components*, Math. Finance, 1:39-55.

[13] Madan, D.B. and Seneta, E. (1990). *The Variance Gamma (V.G.) Model for Share Market Returns*, Journal of Business, 63(4), 511-524.

[14] Schoutens, W. (2003). *Lévy Processes in Finance*. Wiley Series in Probability and Statistics.